

03-31-00

A

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)*(Only for new nonprovisional applications under 37 CFR 1.53(b))*Docket No.
13414 (NECI1083)

Total Pages in this Submission

TO THE ASSISTANT COMMISSIONER FOR PATENTSBox Patent Application
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

METHOD FOR MATCHING A TWO DIMENSIONAL IMAGE TO ONE OF A PLURALITY OF THREE DIMENSIONAL CANDIDATE MODELS CONTAINED IN A DATABASE

and invented by:

A. Peter Blicher, et al.

jc678 U.S. PTO
09/538209
03/30/00If a **CONTINUATION APPLICATION**, check appropriate box and supply the requisite information:☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. ☒ Filing fee as calculated and transmitted as described below
2. ☒ Specification having 52 pages and including the following:
 - a. ☒ Descriptive Title of the Invention
 - b. ☐ Cross References to Related Applications *(if applicable)*
 - c. ☐ Statement Regarding Federally-sponsored Research/Development *(if applicable)*
 - d. ☐ Reference to Microfiche Appendix *(if applicable)*
 - e. ☒ Background of the Invention
 - f. ☒ Brief Summary of the Invention
 - g. ☒ Brief Description of the Drawings *(if drawings filed)*
 - h. ☒ Detailed Description
 - i. ☒ Claim(s) as Classified Below
 - j. ☒ Abstract of the Disclosure

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
13414 (NECI1083)

Total Pages in this Submission

Application Elements (Continued)

3. ☒ Drawing(s) *(when necessary as prescribed by 35 USC 113)*
- a. ☒ Formal Number of Sheets 3
- b. ☐ Informal Number of Sheets _____
4. ☒ Oath or Declaration
- a. ☒ Newly executed *(original or copy)* ☐ Unexecuted
- b. ☐ Copy from a prior application (37 CFR 1.63(d)) *(for continuation/divisional application only)*
- c. ☒ With Power of Attorney ☐ Without Power of Attorney
- d. ☐ DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application,
see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. ☐ Incorporation By Reference *(usable if Box 4b is checked)*
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied
under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby
incorporated by reference therein.
6. ☐ Computer Program in Microfiche *(Appendix)*
7. ☐ Nucleotide and/or Amino Acid Sequence Submission *(if applicable, all must be included)*
- a. ☐ Paper Copy
- b. ☐ Computer Readable Copy *(identical to computer copy)*
- c. ☐ Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. ☒ Assignment Papers *(cover sheet & document(s))*
9. ☐ 37 CFR 3.73(B) Statement *(when there is an assignee)*
10. ☐ English Translation Document *(if applicable)*
11. ☐ Information Disclosure Statement/PTO-1449 ☐ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Acknowledgment postcard
14. ☒ Certificate of Mailing
- ☐ First Class ☒ Express Mail *(Specify Label No.):* EL357933440US

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
13414 (NECI1083)

Total Pages in this Submission

Accompanying Application Parts (Continued)

15. ☐ Certified Copy of Priority Document(s) *(if foreign priority is claimed)*

16. ☐ Additional Enclosures *(please identify below):*


| |
|--|
| |
|--|

Fee Calculation and Transmittal

CLAIMS AS FILED

| For | #Filed | #Allowed | #Extra | | Rate | Fee |
|--|--------|----------|--------|---|---------|------------|
| Total Claims | 28 | - 20 = | 8 | x | \$18.00 | \$144.00 |
| Indep. Claims | 12 | - 3 = | 9 | x | \$78.00 | \$702.00 |
| Multiple Dependent Claims (check if applicable) <input type="checkbox"/> | | | | | | \$0.00 |
| BASIC FEE | | | | | | \$690.00 |
| OTHER FEE (specify purpose) <u>Assignment</u> | | | | | | \$40.00 |
| TOTAL FILING FEE | | | | | | \$1,576.00 |

- ☒ A check in the amount of \$1,576.00 to cover the filing fee is enclosed.
- ☒ The Commissioner is hereby authorized to charge and credit Deposit Account No. 19-1013/SSMP as described below. A duplicate copy of this sheet is enclosed.
- ☐ Charge the amount of _____ as filing fee.
- ☒ Credit any overpayment.
- ☒ Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
- ☐ Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).


Signature

Paul J. Esatto, Jr.
Registration No. 30,749

Scully, Scott, Murphy & Presser
400 Garden City Plaza
Garden City, New York 11530
(516) 742-4343

Dated: March 30, 2000

cc:

METHOD FOR MATCHING A TWO DIMENSIONAL IMAGE TO ONE OF A PLURALITY
OF THREE DIMENSIONAL CANDIDATE MODELS CONTAINED IN A DATABASE

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates generally to computer vision and more particularly to the recognition and identification of a face or other object from among a stored database of three dimensional (3D) models when presented with an arbitrary two dimensional (2D) photograph of a face or object, under arbitrary pose and lighting conditions.

2. Prior Art

There are two subfields of computer vision that are closely related to the field of this invention: face recognition, and object recognition. Face recognition has primarily been concerned with identifying 2D images of human faces, or matching 2D images of faces to other such 2D images. The methods of the prior art have not used 3D models of faces to identify 2D images.

The field of object recognition, however, has had a major concern with identifying 2D images based on known 3D models. There have been numerous works in this area. A more recent example is the methods that use what is called "alignment".

The usual methods consist of extracting some features from the gray scale image which are then matched with candidate features from the models. Examples of features are special points (such as corners), extracted edges, configurations of edges (such as parallel edges, sometimes referred to as "ribbons"), etc. The primary problems in the object recognition literature revolve

around how to (1) choose good features, (2) how to manage the combinatorial explosion that results from the possibilities of matching large numbers of image features with large numbers of model features, and (3) how to measure the quality of a match in the presence of numerous sources of error. The issue of lighting enters primarily at the (early) stage when features for matching are being extracted. The prior art related to object recognition has considerable bearing on performing step (1) above (pose estimation), but has little to say about the problem of accounting for the effects of light that are needed in steps (2) and (3).

Face recognition, on the other hand, has had considerable interest in dealing with lighting. The face recognition literature has evolved a number of methods of dealing with variations in lighting. These revolve around two main approaches: (1) estimating the lighting conditions (as in step (2) above), or (2) analyzing the space of images that is created as the lighting is varied. Typically, method (2) takes advantage of "principal components analysis" (PCA) to identify the dimensions of the image space where the most variation occurs. These typically image-based methods have been fairly successful under conditions where pose does not vary. However, in general they have not been concerned with the use of 3D models, although there are those in the art who have constructed 3D models from images, and have been concerned about the relations between images and 3D structure, as have a large fraction of workers in computer vision.

There is also significant work by others in the art on using the space of distortions of a face in conjunction with the observed shading to permit a linearization that in turn admits a PCA analysis on which recognition can be based. These artisans use a method based on distortions of graphs that join standard feature

points on the face or similar methods, which have been applied in commercially available systems. Many groups in recent years have been interested in face recognition and the literature that has arisen in the past 5 or so years is quite large. The primary focus of all these methods, however, has been on identifying 2D images based on training sets of other 2D images.

A major contributing factor to the explosion in interest in face recognition was the successful application of PCA methods to the problem. This was impressive because it was a tractable solution to what had hitherto been considered an intractable problem. However, there are severe limitations with respect to variation of the conditions under which the data must be captured; in particular, lighting, pose, and scale must be as constant as possible. These limitations can be traced to the fact that PCA is a linear decomposition and therefore only will give good results when the space it is being applied to is a linear space. The space of images under varying pose is not linear, and therefore PCA breaks down. Those in the art have addressed this problem by finding particular ways to convert the nonlinear space back to a linear one that is amenable to PCA. Still others in the art have pointed out that in the simplest cases, any image of such a simple object was just a linear combination of the images under almost any 3 different lighting conditions. This led to activity in this area, culminating with work on rigorously working out the relationship between lighting and the space of images. It should be pointed out, however, that in the presence of shadows, the situation is still quite complex.

There are apparatus known in the art for capturing 3D models of people's faces. These apparatus capture pictures of the faces in registration with the 3D models that are captured. The 3D models consist of the 3-dimensional coordinates of a large number of points on the surface of the face, typically on the order of 640

x 640, along with the color picture value at each point. This provides the possibility of realistic computer graphics rendering of the faces from any vantage point.

5 It is envisioned that these apparatus can be used to capture a large number of face models. In some applications, this number could reach on the order of one million faces or more.

10 One application of such apparatus is face identification, i.e., given a query consisting of an arbitrary 2D photograph of a face, to search a large database of previously captured 3D models of faces to find the model which best matches the photograph.

15 There are numerous obstacles to be overcome to accomplish this. In general, there is no knowledge or control of the conditions under which the query photograph was acquired. Therefore, the pose (i.e. the position and orientation of the subject) and the lighting conditions of the query photograph are unknown.

20 One prior art paradigm, in this setting, for matching a query to a model, proceeds as follows:

25 1) the pose of the face in the query is determined. For purposes of this application, "pose" is defined as the 3D position and orientation of the face;

30 2) the lighting in the query is determined. This means finding the direction, intensity, and color of any number of light sources which may have been illuminating the subject when the query photograph was acquired;

3) for each model in the database, computer graphics techniques are then used to render a realistic image of the model in the pose and lighting conditions determined in steps (1) and (2); and

4) among the renderings computed in the previous step, the one which most closely approximates the query is found.

5 All of these steps involve difficulties which the current state of the art has not overcome. The methods of the present invention concerns the second, third and fourth steps, namely, finding the lighting conditions in the query, and determining which proposed candidate would appear most like the query under those lighting conditions.

The impediment in the method of the prior art described above is that to solve for lighting requires knowing the 3-dimensional configuration of the surface that was photographed, as well as its reflectance characteristics. However, the query does not provide any of this information directly. A large effort in computer vision over the past 40 years has been devoted to inferring this information from one or more pictures.

The methods of the prior art perform steps (2) and (3) above by first solving for lighting by solving the image irradiance equation (the Lambertian imaging equation), which is a system of linear equations, with 5 independent variables for each potential light source: 2 variables indicating the direction of the light (for example, bearing and azimuth), and 3 variables indicating the red, blue, and green intensity, respectively. The number of equations is equal to the number of data points in the image for which surface normal and reflectance are available times the number of color components, typically 3. However, an important constraint is the requirement that the light intensities must not be negative. This latter requirement precludes solving the system using a standard method for linear systems, and instead requires a more sophisticated method, such as a linear programming technique, or the algorithm known as nonnegative

least squares. These special techniques require significantly greater computer time to solve than the same system would require without the constraints of non-negativity. This approach was brought forth by those in the art, as mentioned above.

5

SUMMARY OF THE INVENTION

10

Therefore it is an object of the present invention to provide a method for matching a 2D image to one of a plurality of 3D candidate models contained in a database in which an object is identified without the need for solving for light sources or other complex and costly computations such as eigenvector determination.

15

It is a further object of the present invention to provide a method for matching a 2D image to one of a plurality of 3D candidate models contained in a database which has relatively lower costs over prior art methods that solve for lighting.

20

It is yet a further object of the present invention to provide a method for matching a 2D image to one of a plurality of 3D candidate models contained in a database which is simpler than prior methods.

25

It is still yet a further object of the present invention to provide a method for matching a 2D image to one of a plurality of 3D candidate models contained in a database which does not require iterative algorithms that take one to two orders of magnitude more computer time on the size of problems involved.

30

Accordingly, a method for matching a two dimensional image to one of a plurality of three dimensional candidate models is provided. The method comprises the steps of: determining the position and orientation of an object giving rise to the two dimensional

image; computing a representation of lighting effects that allows the lighting that gave rise to the two dimensional image to be used to render a realistic image of a three dimensional model; successively rendering each three dimensional candidate model in the determined position and orientation using the surface normals in conjunction with the corresponding computed representation of lighting effects; and comparing the two dimensional image with each of the rendered three dimensional candidate models.

Another method for matching a two dimensional image to one of a plurality of three dimensional candidate models is also provided. This method comprises the steps of: determining the position and orientation of an object giving rise to the two dimensional image; computing a representation of lighting effects based on the plurality of three dimensional candidate models and the two dimensional image that allows evaluation of the likelihood that a particular three dimensional candidate model gave rise to a particular two dimensional image; and choosing the most likely model to have generated the query based on the computed representation of lighting effects.

In a first variation of the methods for matching a two dimensional image to one of a plurality of three dimensional candidate models, the method comprises the steps of: determining the position and orientation of the two dimensional image; for each three dimensional model, computing a histogram-like table having a computed brightness coefficient for each surface normal of the model based on the corresponding value in the two dimensional image, each brightness coefficient depending only on the corresponding surface normal; successively rendering each three dimensional model in the determined position and orientation using the surface normals in conjunction with the corresponding computed brightness histogram-like table; and comparing the two dimensional image with each of the rendered

three dimensional models.

In a second variation of the methods for matching a two dimensional image to one of a plurality of three dimensional candidate models, the method comprises the steps of: determining the position and orientation of the two dimensional image; for each three dimensional model, computing a histogram-like table having a computed brightness coefficient for each surface normal of the model based on the corresponding value in the two dimensional image, each brightness coefficient depending only on the corresponding surface normal; computing the variance of the brightness coefficients that are used to create each bucket of the histogram-like table, a bucket being a set of similar normal values that are lumped together to form a single argument value in the histogram-like table; computing the sum of the variances of the histogram-like table buckets; and ranking the three dimensional models using its computed sum as an error function, the ranking indicating the likelihood that the corresponding three dimensional model matches the two dimensional image.

Still further provided is a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of the methods and variations thereof of the present invention and a computer program product embodied in a computer-readable medium for carrying out the methods, and variations thereof, of the present invention. Preferably, the computer program has modules corresponding to the steps of the methods and variations thereof of the present invention.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other features, aspects, and advantages of the methods of the present invention will become better understood with

regard to the following description, appended claims, and accompanying drawings where:

FIG. 1 illustrates a schematical representation of the method of the present invention.

FIG. 2 further illustrates the lightsphere module of FIG. 1.

FIG. 3 illustrates a graphical representation of the lightsphere data structure of FIG. 2.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Although this invention is applicable to numerous and various types of objects to be recognized, it has been found particularly useful in the environment of human face recognition. Therefore, without limiting the applicability of the invention to human face recognition, the invention will be described in such environment.

The methods of the present invention, have an efficient method of comparing a query photograph, such as a face of a person to be identified to one of a plurality of 3D models contained in a database. The methods of the present invention presupposes a particular model of reflectance, known as the Lambertian assumption. Although the invention is robust against deviations of the true details of reflectance from this model, the relevant point is that this is a functional relationship between the data in the query and the data in the model, i.e., there is an equation that relates the query and the model. The methods of the present invention can be viewed as an efficient way of measuring to what degree the expected form of this functional relationship holds between the query and the putative model from which it arose. This measure can be generalized; for example, the statistical notion of correlation is a generalization of

functional relation.

Another generalization, which arose primarily in information theory, depends on the entropies of the distributions governing the objects we want to compare. These distributions, for example, could govern the point values of different measurements in the two objects. Entropy can be regarded as a measure of uncertainty, i.e., the greater the entropy, the more difficult it is to predict the outcome of drawing a sample from a probability distribution. When we have two random variables with a joint distribution, e.g. the data values of a 3D model and of a 2D image putatively of that model, it can be determined to what extent does knowing the value of one of these random variables, say a 9-point patch on the 3D model, help us predict the value of the other variable, say the intensity at a corresponding pixel in the 2D image. This question is quantified by taking the difference between the entropy of the marginal distribution of the second variable, and subtracting off the entropy of the conditional distribution of the second variable given the first. This yields how much our uncertainty decreased thanks to knowledge of the first variable. This can be summarized by:

$$I(Y;X) = H(Y) - H(X|Y) \quad (1)$$

where H denotes entropy, X is the first random variable, and Y is the second random variable. This quantity, denoted by $I(Y;X)$ is known as the mutual information between X and Y . It is symmetric, and is also equal to

$$H(X) + H(Y) - H(X,Y) \quad (2)$$

where $H(X,Y)$ is the entropy of the joint distribution of X and Y .

Thus, the methods of the present invention can be thought of as a

particular approximation to measuring the mutual information between the (distributions governing the) 3D model and the 2D query.

5 In a system where a 2D photographic query is to be matched to a 3D model in a large database of such models by predicting the appearance of each model under the pose and lighting conditions that created the query, the methods of the present invention avoid the problem of solving for the light sources that produced
10 the query image by using each successive 3D model as a candidate solution, and in conjunction with the query, building a histogram-like table that specifies a brightness coefficient for each surface normal which depends only on the normal. In one variation of the methods of the present invention, the model is then rendered in the already approximately known pose using these brightness coefficients, and the resulting image is compared to the query image and an error is computed. This error determines a ranking of the models for the purpose of identification. In another variation, no rendering is done, and instead the variance of the samples in each "bucket" of the brightness vs. normal
20 histogram-like table is computed, and this variance, possibly weighted by the number of samples is summed over all buckets to yield a quality measure, which is then used for identification in the same way as an error measure.

25 It is assumed that there is a good, though not necessarily exact, solution for the pose of the query. However, it is not known which model generated the query image, and it is not known what were the lighting conditions for the query.

30 The method of the present invention therefore proceeds as follows:

a. Assume, in turn, that each model is the "correct" one.

b. Using the query image, compute a histogram-like table that associates a "brightness" coefficient with each normal. (For color images, three such coefficients are preferably computed, one for each color component, typically red, green, and blue).

5

There are then 2 preferable variations for the method. The first is as follows:

10

c1. Use the normals of the model under consideration in conjunction with the normal/brightness histogram computed in step (b) to render a picture of the model under putatively the same lighting conditions as were present when the query was photographed.

d1. Compare the query with the rendered picture to rank the likelihood that this model is in fact the correct one.

The second variation replaces the last 2 steps with these:

20

c2. Having saved additional information when constructing the normal/brightness histogram, compute the variance of the brightness coefficients that were used to compute each bucket of the histogram. (A "bucket" is the set of similar normal values that are lumped together to form a single argument value in the histogram.)

25

30

d2. Compute the sum of the variances of the histogram buckets, weighted by the number of samples in each bucket. Use this sum as an error measure in ranking the models to select the correct one.

In step (b), solving for the lighting in the paradigm that was discussed above as step (2) is avoided. However, in step (c1), it is still possible to perform the rendering that is called for

in step (3) of the paradigm. In the alternate evaluation steps, (c2) and (c3), it is possible to avoid rendering altogether in assessing the likelihood of a match between a 3D model and a 2D photograph query.

5

The methods of the present invention are thus based on the observation that computing the light sources, i.e. the lighting, that produced the query image, is only an intermediate step in a process whose real goal, in the context of the overall system, is to render the identification candidate under those lighting conditions for the purposes of comparison to the query, or indeed to compute the relative likelihood that the model matches the query even without doing such a rendering but with equivalent results.

10

What is ultimately being measured is the consistency of the imputed brightness coefficients associated with each normal. Because the 3D model's normals and reflectance values are used in conjunction with the query's observed intensities, this is a measure of the error in fitting the 2D photograph query to the 3D model.

20

The methods of the present invention are based on a Lambertian model of reflectance. The Lambertian assumption is that the intensity projected onto a point in a monochrome photograph of an object lit by a point source of light depends only on the following quantities:

25

R, the "reflectance" at the point of the object that gives rise to the corresponding point in the picture,

30

N, the "normal" to the surface at the point in question (i.e. the unique perpendicular direction to the surface at that point),

L, a unit vector representing the direction from the point to the light, and

b, the intensity of the light source at the point.

5

In particular, the observed intensity in the picture does not depend on the viewpoint, but only on the intrinsic properties of the object, and the lighting direction and intensity.

10 The observed image intensity is described by the equation:

$$I = R \cdot b \cdot L \cdot N \quad (3)$$

where I is the image intensity at the image point in question, and \cdot is the usual vector dot product.

If the quantities R, b, L, and N are given at each point, then in order to make a rendering from a 3D model, one need only perform the multiplications in equation 3. Also, it must be determined which portions of the model will be visible; however, that can be performed adequately by existing computer graphics techniques.

The situation in the model identification problem being solved is that R and N are known for each model, I is known from the query, and b and L are unknown. One approach, involving step (2) in the paradigm discussed above, is to solve for b and L, and then perform the multiplication in equation 3.

The approach of the methods of the present invention is to attempt to solve only for the composite quantity $b \cdot L \cdot N$, and then (in the first variation mentioned above) perform the multiplication in equation 3.

At each point, equation 3 is divided to give

$$b L * N = I/R \quad (4)$$

call this quantity $B(N)$, a brightness coefficient.

5 After rendering, equation 3 becomes

$$I_{\text{rendered}} = B(N) R \quad (5)$$

10 If this were done at just one point, this is a triviality, and we would simply get back the query's 'I' that we started with. However, in general, for every normal that occurs in the model, there will be many other points with the same or nearly the same normal. If all the B values that we get are averaged whenever the same normal occurs, and equation 5 is used to render a picture from the model, we will get only an approximation to the 'I' in the query. If the different values that were averaged to obtain $B(N)$ for a given N were different due to small variations in material properties, measurement error, nonlinearities, etc., but the query did in truth match the model, then it can be expected that the average represented by $B(N)$ would be a good approximation to the number needed for equation 5 to give a good approximation to the observed query intensity. However, when the query does not in truth match the model, because the values of N that are being used come from the model, while the intensities, I, come from the query, it can be expected that the correlation among $B(N)$ values will be poor, and the average, when inserted into equation 5 for rendering, is unlikely to give a result which is close to the value in the query. In that case the rendering will be found to have a larger error from the query than in the case where the query in truth did match the model.

15
20
25
30

B is a function which takes a surface normal vector to a number, i.e., $B: S^2 \rightarrow R$ where S^2 is the so-called Gaussian sphere, which is the space (or manifold) of all possible normal

directions to a surface embedded in 3 dimensional space. The 2
denotes the fact that this sphere is 2-dimensional, like the
surface of a common ball. For this reason, the method of the
present invention is referred to as the "lightsphere" method,
5 since it attempts to find a brightness function defined on the
Gaussian sphere.

So far, the situation has been presented with a single point
light source and a monochrome picture. To generalize from
10 monochrome to color, the monochrome case is done 3 times, once
for each color component.

Generalizing the lighting, other than its color, is not as
simple. In general the lighting can be, and usually is, much
more complex than a simple point source. Any point of the
surface can have light impinging on it from all directions, in
the hemisphere visible to it, and this lighting can vary from
point to point. At each point, the lighting is also a function
on the Gaussian sphere, prescribing the intensity of light coming
from the direction represented by that point on the Gaussian
sphere. In general, this function can vary from point to point
throughout space, and therefore can vary on the surface we are
imaging. If the further assumption is made that all the lights
are infinitely far away, then the lighting function does not
25 vary from point to point, and the global lighting can be thought
of as a function on the Gaussian sphere.

To generalize equation 3 to this situation, we simply integrate
over the Gaussian sphere, or approximate the integration by a
30 summation. There is a small complication in that a nonlinearity
is introduced because only the lights with a direction that
result in a positive dot product with N make a contribution;
there is no negative brightness from lights behind the surface.
However, this does not disturb the key point, which is that the

new lighting equation, which is now something of the form

$$I = R f(N) \quad (6)$$

5 still uses the exact same light distribution whenever the same
normal is encountered, and therefore the computation of equation
4 and the rendering of equation 5 are just as valid as with a
single point source. It should be noted that for both equation 3
and equation 6 to be independent of position and depend only on
10 normal, it is assumed that there are no "cast shadows". Cast
shadows are shadows that are caused by non-local interactions,
such as when the nose casts a shadow onto the cheek; these are
distinguished from "attached shadows" which occur as the object
smoothly curves away from the light source, as occurs when only
the lit part of the moon is visible at night. The methods of the
present invention neglect the effect of cast shadows.

15 In fact, with more complex lighting, when we are attempting to
match a wrong model, the fact that different light sources will
come into play in the query, for the same normals on the model,
will result in an even larger error when matching the wrong
model, with no similar increase in error when matching the
correct model. Hence, this lighting nonlinearity actually
enhances the effectiveness of the invention.

20 A similar observation applies with regard to specularities. The
location of specularities in an image violates the Lambertian
assumption of equation 3 in two important ways:

- 25
- 30 1. The location of the specularity depends on the camera
viewpoint; and
 2. The brightness of the specularity does not vary linearly with
either the lighting vector or the normal vector.

These problems make it difficult to incorporate specularity into any system which must solve for light sources, though it is by no means impossible. Consequently, specularities often lead to increased error.

However, if we continue to assume that the light sources are at infinity, and there are no cast shadows, then even in the presence of specular reflection, equation 6 still applies approximately. More specifically, the new equation is

$$I = R f(N) + S g(N) \quad (7)$$

where $S g(N)$ represents the specular component of the reflected light. S is the shininess of a particular point of the surface, in analogy to R , while $g(N)$ captures the intensity variation as a function of the normal.

When we apply the lightsphere calculation of equation 4 to equation 7 we get

$$I/R = f(N) + S g(N)/R \quad (8)$$

However, if this is to be $B(N)$, a problem arises because R and S depend on the position on the surface, and not on the normal. This means that the quality of the rendering that can be achieved with the $B(N)$ that can be computed as the average of the above values I/R will depend on the degree of variation of R and S at the normals involved. Nevertheless, the rendering will approximate some specular component where it exists, and will therefore be closer to the exactly correct rendering than one based simply on rendering from computed lights with a Lambertian rendering equation. Of course, with another method, one can compute the lights based on some nonlinear model that includes specularity, and render with that model. However, in the methods

of the present invention, this is achieved at no increase in complexity; with other methods, the specular reflectance equation would have to be explicitly modeled, with a concomitant increase in complexity and computational cost.

5

Conversely, when the wrong model is rendered for the query, because specularities tend to have quite bright components, wildly inappropriate values for $B(N)$ are expected when there are specularities. This leads to proportionately greater error added to the wrong model's ranking than to the right one's. Hence, specularity in the query can actually enhance the effectiveness of the invention.

10

We note that the measures corresponding to the two variations of the methods of the present invention are closely related. The two measures are, respectively, (1) error between rendered image and query, and (2) weighted variance of lightsphere coefficient values within each bin. Each coefficient that entered into the variance computation arose from a point in the query. When that point is rendered with the average lightsphere coefficient for its normal's bin, the error will be proportional (via the reflectance) to the difference between the average coefficient being used in the rendering and the exact coefficient that was used in computing the average. Were it not for the reflectance multiplier, when summing over the points that fall into the bin, the sum of the squares of the one set of errors is the same as the sum of squares of the differences, and the latter is exactly the variance of the bin weighted by the number of samples. Thus, the two measures differ only by a weighting factor of reflectance in the error sums. It is to be noted that by error between images, it is meant RMS (root mean square) error, or equivalently, SSD (sum of squared differences). Other error measures are possible, as well.

20

25

30

under the Lambertian model we expect to be the same as
b L * N (2). The value of N (220) that is known for this point
from the model 10 is the value used to choose the correct bin 200
of the data structure 160, and the value I/R is then inserted
5 (230) into this bin 200, either by summing, or recording the
values inserted into the bin 200, depending on which variation of
the methods of the present invention is being used.

Although the model evaluation process has been presented as a
10 sequential one, where only one model is being evaluated at any
time, as discussed above, it is possible to divide attention
among numerous models using partial solutions so that the cost of
fully evaluating a model is only expended if the model continues
to have a low error in competition with other models.

Also, prior to the final error ranking process, lightsphere
results for one model do not require results for any other model.
Consequently, the overall process trivially admits of massive
parallelization, with a proportional increase in overall speed.

The lightsphere module 60 will now be described in greater
detail. The lightsphere module 60 can be divided into two
stages:

25 1. Preprocessing steps that:

a. Compute the pose of the model that best accounts for the
query.

30 b. Generate a registered map of the surface normals of the
model corresponding to points on the query.

c. Generate a registered image of the model's reflectance
corresponding to points on the query.

2. Using the information provided in stage 1, in conjunction with the query image, to compute the lightness coefficients on the Gaussian sphere.

In stage 1, the pose is computed by any of numerous methods of the art which can be used for this purpose. Generating the registered maps of normals and reflectance is done by standard and well-known computer graphics methods. Note that computing the reflectance that is associated with the 3D model is not a simple procedure, but since such procedures are known in the art, a description of its computation is omitted. One such method for computing the reflectance associated with the 3D model is to simply use images taken under flat lighting conditions.

The input to the preprocessing stage consists of a query image Q , a candidate 3D model M , and a camera model K . The output of the preprocessing stage comprises a pose T , a registered reflectance image R , and a registered normal map N . The input to the lightsphere module consists of the query image Q , the registered reflectance image R , and the registered normal map N . The output of the lightsphere module comprises a normal-lightness table B , and an error measure E .

These inputs and outputs will now be precisely explained, as will exactly how to compute the lightsphere output. Note that the definitions being used in this section differ slightly from the definitions used above.

For simplicity, only a monochrome image is considered. For a color image, the process to be described below is repeated separately for each color component, e.g., red, green, and blue.

The query image Q is considered to be a function defined on the real plane, i.e.,

$$Q: R^2 \rightarrow R^+ \quad (9)$$

This notation means that Q is a map from the 2-dimensional real Euclidean plane to the nonnegative real numbers. It is assumed that Q only takes nonzero values, or is only defined, in some finite rectangle of the plane. In actual fact, Q is sampled at only a finite discrete grid of points, so that it takes values at only integer coordinate values. Note that Q is a function, while the notation $Q(p)$, where p is a point in R^2 , denotes the real number to which Q maps the point p .

The 3D model M consists of several subparts. These include a set of n points

$$\{p_i \mid i = (1, 2, \dots, n)\} \quad (10)$$

in R^3 which are the locations of sample points on the model, and a corresponding set of n real numbers r_i which represent the surface reflectance at each corresponding p_i .

The camera model K is a perspective projection. I.e.,

$$K: R^3 \rightarrow R^2 \quad (11)$$

where, after some change of coordinate systems by rigid motion,

$$K(x,y,z) = a^*(x,y)/z, \quad (12)$$

for some positive real number a . The change of coordinates is part of the camera model. There are many other ways to formulate the definition of perspective projection; this one is equivalent, and the exact formulation is not important. For the purposes of the lightsphere computation, distortion effects of the camera are

not considered, and the camera model is fixed. Variations where more information about the camera is used are obvious. Note that this definition allows for some amount of perspective distortion, depending on the values of the camera model K and the position of the 3D model in space. However, to the degree that the true camera and subject location for the query lead to more or less perspective distortion, this difference is not compensated for in what is described below, although a more detailed camera model could accomplish this.

The pose T is the rigid motion in R^3 which, when applied to the model M , after projection through the camera model K , brings the image of the model M into the best registration with the query Q .

I.e., the pose T tries to satisfy the following equation:

$$K(T(p_i)) = q_i \quad (13)$$

where p_i is a point of the 3D model, and q_i is the corresponding point of the query image, simultaneously for some set of points p_i and q_i .

The registered reflectance image R is a map

$$R: R^2 \rightarrow R^+ \quad (14)$$

(Note that the first R is a different R than the 2nd and 3rd.) sampled on the same grid as the query image Q , such that

$$R(K(T(p_i))) = r_i \quad (15)$$

In other words, R tells us, given a point in the query image Q , what the reflectance is of the point in the model that projects to that point in the query given the pose and camera model. Note

that several points on the model might project to a single point in the query; i.e. the computer graphics technique of hidden surface elimination is necessary to find the unique R that corresponds to the surfaces of the model that are actually visible under the assumed pose and camera model. Numerous well-known computer graphics techniques exist that can accomplish this. Both proprietary and freely available packages have been used to do this computation with equal success.

The surface normal n_i to a point p_i of the 3D model M is a unit vector in R^3 which represents the direction at right angles to the surface of M at p_i . Computing n_i requires knowing some nearby points of M, i.e. nearby on the surface that M represents. In our apparatus, this is known in advance because the data in M comes from a system that essentially produces a raster scan. In addition, there are standard computer graphics techniques that allow putting a less ordered set of points into a form where such neighboring points are known, for example, by the process known as triangularization of the surface. If p has the neighbors q and s, then the unit normal at p can be computed by

$$v = (p-q) \times (p-s) \quad (16)$$

$$n = v / ||v|| \quad (17)$$

where \times is the vector cross-product, and $||v||$ is the norm of v.

In practice, it is often better to consider several neighbors of p, e.g. in a grid, the 8 or 4 neighbors, and take the average of the values thus computed.

Because n is a vector in R^3 it can be written as

$$n = (n_x, n_y, n_z) \quad (18)$$

where the three components are the components of n in the x , y , and z directions, respectively. However, the requirement that n be a unit vector imposes the constraint that

$$1 = ||n||^2 = n_x^2 + n_y^2 + n_z^2 \quad (19)$$

This can be viewed as the equation of a 2-sphere, i.e., the surface of a ball in 3-space. Geometrically, if n is considered to be anchored at the origin, as n varies among all possible values for a unit vector, its tip varies over the surface of the unit sphere. Therefore each n can be identified with these tip positions, i.e. with the points of the 2-sphere. The 2-sphere is commonly denoted by S^2 .

When this identification is made between the unit normals of a surface and the 2-sphere in differential geometry, the sphere is commonly referred to as the "Gaussian sphere", and the map which maps each point of the surface to its normal is known as the "Gauss map".

With this in mind, the registered normal map can be defined as a map

$$N: R^2 \rightarrow S^2 \quad (20)$$

which takes each point in the query image Q , to the surface normal at the point of the 3D model M which projects to that point of Q under the pose T and camera model K . I.e., N is defined to satisfy

$$N(K(T(p_i))) = n_i \quad (21)$$

in analogy to the similar equation (15) for the registered reflectance, where now n_i is the normal calculated in the

preprocessing stage at the point p_i .

15 N, again, is calculated only at the grid points associated with
the grid points of Q. The calculation of N, like the calculation
5 of R, requires hidden surface removal to guarantee that only the
value of n for the visible surface is used. Again, this is most
easily done by using existing standard computer graphics
techniques. One can take advantage of existing graphics
10 interfaces to do this by allowing the graphics engine to compute
renderings with 6 different light sources: one for each
combination of positive and negative coordinate direction, i.e.,
lights coming from +-x, +-y, +-z. With a Lambertian rendering
model in the graphics engine, these will yield the positive and
negative components of the normal vector at each point. This
15 also allows the use of the graphics engine's techniques for
interpolation, which helps considerably in the performance of the
algorithm.

20 The lightsphere algorithm calculates a normal-lightness table B
from the query Q, the registered reflectance image R, and the
registered normal map N.

25 Before we define and describe how to compute B, consider first
the map

$$A: S^2 \rightarrow R^+ \quad (21)$$

given by

30
$$A(n) = Q(q)/R(q) \quad (22)$$

when $N(q) = n$.

For notational compactness, define

$$D: R^2 \rightarrow R^+ \quad (23)$$

$$\text{by } D(q) = Q(q)/R(q). \quad (24)$$

5 (Elsewhere this is referred to as I/R.)

We would like A to be defined by $A(n) = D(N^{-1}(n))$. However, there may be many points where the normal is n, so that $N^{-1}(n)$ is not a single point, but a set of points. To estimate A, it is
10 desirable to take the average of D on this set of points. I.e., suppose that

$$P(n) = \{p \mid N(p) = n\} \quad (25)$$

is the set of points in the query that map to the normal n. Our goal is to define A by

$$A(n) = \frac{1}{\#(P(n))} \sum_{p \in P(n)} D(p) \quad (26)$$

where $\#(P(n))$ is the number of points in the set $P(n)$ (or the measure of the set, in the continuous case). This is a well-defined quantity.

However, this is not adequate when a discrete grid, in this case the discrete grid of the query image is being operated on. The problem is that the normals are derived from the 3D model, and in
30 general the values to which they have been discretized depend only on the accuracy of the 3D model data and the method whereby the normals were computed. In particular, they will generally be discretized to such a fine degree that it is unlikely that any two of the discretized points q_i of the query will map to the
35 same normal n, and therefore almost every $A(n)$ will arise from only a single sample. In that case, the lightsphere calculation,

when used to render an image from the model will be trivial,
i.e., it will result in exactly the query image value, since we
will simply be computing the quantity

5
$$I(p) = A(N(p)) \quad R(p) = D(p) \quad R(p) = Q(p)/R(p) \quad R(p) = Q(p) \quad (27)$$

10 Therefore, it is essential that a way be devised that in general,
several values will be averaged to obtain $A(n)$. A way to do this
is to discretize the Gaussian sphere into a small enough number
of bins (or buckets) that a sufficient number of buckets contain
more than one sample. There are numerous ways to discretize the
2-sphere. For example, one can take an icosahedron centered at
the same point as the sphere and radially project the icosahedron
to the sphere. This gives a partition, or tessellation, of the
sphere. However, for our purposes it has been found that it
works well to use an extremely simple method of discretizing the
sphere, which can now be described.

We start with a unit normal vector

20
$$n = (n_x, n_y, n_z) \quad (28)$$

25 In the methods of the present invention, surface normals that
point away from the camera are not of interest, because the
camera cannot see any surface point that has such a normal. By
"point away from the camera", it is meant that if v is a vector
that points in the direction of the camera perspective projection
center from the point p on the surface, and n is the outward
pointing normal at the point p , then if $n * v < 0$, where $*$ is the
30 vector dot product, we say that n is pointing away from the
camera. Thus only normals n on half of the Gaussian sphere are
of interest, the half where the above inequality does not hold.
Let us for the moment use a coordinate system where the camera
axis lies along the z axis and the camera is in the positive z

direction from the subject. Then the set of all surface normals we are interested in is exactly those where $n_z > 0$. This can be thought of as a hemispheric cap lying on the x-y plane.

Now our goal is to discretize this hemispheric cap based on the values n_x , n_y , n_z . The easiest way to do this is to ignore the z component altogether, and consider only the projection of the hemispheric cap onto the x-y plane, which is a circular disk. We will discretize this disk, and define the discretization of the hemisphere by the projection back up along z from the plane to the hemisphere. To discretize this disk, the unit disk in the plane, we will simply discretize the unit square in the plane, and restrict it to the disk. To simplify the next discussion, consider the unit square to be the part of the plane given by

$$0 \leq x < 1 \quad \text{and} \quad 0 \leq y < 1. \quad (29)$$

(Since the values of the unit normal components vary from -1 to +1, this involves a scale and shift; in an actual program, one would probably not bother to do this).

The way to do this is as follows. We will discretize so that there are m bins, i.e., into an m x m square tessellation. Divide the unit square into m^2 subsquares S_{ij} , $i, j = 1, 2, \dots, m$, as follows. A point p of the unit square is in subsquare S_{ij} if

$$(i-1)/m \leq p_x < i/m \quad \text{and} \quad (30)$$

$$(j-1)/m \leq p_y < j/m \quad (31)$$

where p_x is the x coordinate of p and p_y the y coordinate.

Then considering each subsquare S_{ij} a "bin" or "bucket", we say that p falls into the (i,j)th bucket if the above conditions (30)

and (31) are satisfied.

Using this definition, a unique bucket (i,j) can be defined for each unit normal n as follows. First, because a shift and scale was done, define

$$n' = [n + (1,1,0)]/2 \quad (32)$$

This moves and shrinks the hemisphere to lie over the unit square as it was defined above. (This is only necessary because simple notation is desirable; this entire analysis is identical with a square that goes from -1 to +1 in x and y .)

Now just compute i and j according to (30) and (31), using n'_x in place of p_x and n'_y in place of p_y , n_z is simply ignored, which has the effect of projecting the square up to the hemisphere. These (i,j) coordinates define the bins for the basic lightsphere procedure. Note that these are not of uniform size on the Gaussian sphere; however, this is not very significant in practice, as they are uniform in the query image.

Now the algorithm can proceed by systematically iterating through the points of the query image. First, it must keep 2 items of data for each bin: a running sum, and a count of the number of samples used thus far in the sum. (Since the count is used to compute an average, an updating method could be used to maintain only a single number for each bin, but this is not worth the bother since the table is small.) Then as the algorithm iterates through the query image Q , at each point p it refers to the reflectance image R and the normal map N . It computes the value of $D(p) = Q(p)/R(p)$ and adds this value to the (i,j) th bin where (i,j) are computed according to (30) and (31) from the value of the normal $N(p)$. (With the understanding that by this it is meant that the same scale and shift operation that gave n' above

is taken into account, here and in what follows.) Also it increments the count for the (i,j)th bin. Once it has finished iterating through the query image, it divides the sum at each bin by the count at that bin to arrive at an average. This results in a table of numbers B(i,j) that are the B(n) mentioned above, the lightness table.

In practice, as explained above, additional criteria are used to decide whether to include a particular query image point in the above computation. For example, image points that are too bright or too dark are excluded.

So far the methods of the present invention have been explained for the variation where the error between a rendered image of the model and the query image as our quality criterion is used.

To render an image using the lightsphere data structure B(i,j) the algorithm proceeds as follows. It iterates through the registered reflectance image R. At each point p of R, it looks up the normal N(p) in the registered normal map N. Using the value of N(p) thus retrieved, it computes the corresponding lightsphere bin (i,j) using (30) and (31), and retrieves the value B(i,j). It then computes

$$I(p) = R(p) B(i,j) \quad (33)$$

and sets the brightness value of the rendered image I to that value, after clipping to avoid values larger than the maximum brightness.

Actually, there is an additional step before the last computation is done. Because there are a very small number of bins compared to the number of actual normal values, using the value B(i,j) directly as above results in considerable artifact due to the

coarse digitization of the bins. Therefore, instead an interpolated value $B'(p)$ is used. The interpolation is based on the value of the normal $N(p)$. Since we know where $N(p)$ falls in the square defined by the (i,j) th bin, we simply use bilinear interpolation with the adjacent values of $B(i,j)$ to arrive at the interpolated $B'(p)$, which is then used to compute

$$I(p) = R(p) B'(p) . \quad (34)$$

The error E is then computed as the sum of square differences or root mean square error between the query image Q and the rendered synthetic image I , which are already registered by virtue of the method of construction of I .

Now we return to the variations of the algorithm that use the variances of the lightsphere. This proceeds identically to what was described above, except that the data that is associated with each bin is slightly more complex. Rather than keep a running sum and a count for each bin, instead for each bin a list is kept of all the $D(p)$ values that fell into that bin. When the iteration is completed, the mean and variance of the list of values in each bin is computed for that bin. The mean is identical to $B(i,j)$ defined above. In addition, we also now have the variances $V(i,j)$.

In the first variation using variances, the error E is defined as the sum of these variances, i.e.,

$$E = \sum_{i,j} V(i,j) \quad (35)$$

In the second such variation, we want to weight this sum by the counts, which will be more similar to the sum of square differences between the rendered and query images used earlier.

Call the counts that were kept in the accumulation process
C(i,j). Then in this variation, the error measure is given by

$$E = \sum_{i,j} C(i,j) V(i,j) \quad (36)$$

It should be noted that although a few very specific methods to
accomplish the lightsphere computation have been presented, there
are other formulas and algorithms that could be used for the
various components and for the error term. These equations and
explanations are merely the description of one possible way to
achieve an embodiment of the invention, and many other ways are
possible with simple modifications and/or additions.

EXAMPLE

The methods of the present invention have been actually reduced
to practice in both the C and C++ programming languages and
tested on a pilot database containing 37 3D models, and actual
queries consisting of 37 subjects in all combinations of 7 poses
and 6 lighting conditions.

Several additional techniques can be utilized to optimize the
performance of the methods of the present invention which have
been actually reduced to practice.

1. It has been found that a grid of 32x32 bins for the
lightsphere histogram works well for images and models with a
resolution of 640x640. 24x24 bins also works well. The number
of bins imposes a tradeoff between accuracy in the brightness
coefficient -- the accuracy increases as there are more bins,
since a smaller range of normals is used in the averaging -- and
statistical significance of the average, which improves as the
number of bins is decreased, since a greater number of normals

then fall into each bin. Thus, the methods of the present invention can be further improved by estimating the magnitude of these effects upon seeing the data and choosing an appropriate number of bins.

5

2. It is possible to use a tessellation of the Gaussian sphere for the bins. This improves the performance at the expense of a slightly greater cost in computing the lightsphere histogram and rendering from it. Furthermore, it has been found that
10 projecting a tessellation of the unit disk onto the Gaussian sphere, i.e. using only x and y values, not z, for determining the bins, is quite effective, despite the sampling inequities it introduces.

3. Because the histogram bins are quite coarse relative to the image and model resolutions, it has been found that bilinear interpolation of the lightsphere values when doing rendering greatly improves both appearance and accuracy.

20

25

30

4. For a Lambertian object, the dot product in the integral which gives the image intensity reduces to a cosine between normal and light directions. The resulting image intensity as a function of normal is essentially a convolution of this cosine with the light intensity distribution. This amounts to a low pass filter operation in the absence of cast shadow. Thus, in a system which solves for lighting and then does rendering, there would be enforced a smoothness upon the brightnesses as a function of normal by virtue of this convolution. In the methods of the present invention, the lightsphere computation does not
contain such a mechanism to enforce smoothness. Consequently, the lightsphere data can contain variations with high frequency components that result from multiple sources, such as departure from Lambertian, reflectance errors, etc. Therefore, it has been found useful to do a small amount of smoothing on the lightsphere

histogram. This can be done by a convolution with a simple small box function. This significantly improves the rendering and the performance of the method.

5 5. When using equation 4 to compute the lightsphere data structure, points with very low reflectance can contribute enormous numbers when the division is performed. Such points lead to meaningless results. To avoid this problem, a minimum threshold is used which the reflectance must exceed before it can enter into the computation. This threshold is currently set ad hoc; however, it could also be computed based on inspecting the reflectance histogram, and stored once and for all with each model. Since the computation of this threshold can be performed at the time the model is added to the database, arbitrarily sophisticated methods could be used to compute it. As an extreme example, it could be set to provide optimum differentiation from other models in the database, based on synthetic renderings of numerous models.

20 6. Extreme specularities tend to be associated with points, such as areas in the eye, where the range data is unreliable. In addition, such points tend to saturate the intensity scale. Therefore it is preferable to impose a maximum threshold which intensity values must not exceed in order to be included in the lightsphere computation. Similarly, extremely dark areas of the query are strongly associated either with areas in cast shadow, or areas where reflectance computations mean little (such as hair). Both of these are sources of error for use, thus, a minimum threshold on intensity is also used in the query.

30 Those of skill in the art will appreciate that the methods of the present invention yield an identification of the object without the need for solving for light sources or other complex and costly computations such as eigenvector determination. Thus, the

methods of the present invention result in the following advantages over prior art methods:

1. A primary advantage over prior art methods is an enormous savings in computational cost. This savings can make the difference between practicality and impracticality in a complete system. For example, the methods of the prior art which use mutual information, while more general and perhaps more effective in principle, are prohibitively expensive for an application requiring a large number of comparisons with models.

2. The methods of the present invention are simpler than the prior art methods. The methods of the present invention required mainly the single process of building a histogram-like data structure, which is little more than visiting each point of the model, extracting the pre-stored normal and reflectance, and performing a few arithmetic steps to update the histogram (while the data structure is referred to as a histogram, it is really a slight generalization). By comparison, other methods require iterative algorithms that take one to two orders of magnitude more computer time on the size of problems involved. For example, one prior art method involves use of the so-called Nonnegative Least Squares algorithm, which is a complex iterative algorithm. Similarly, other approaches of the prior art require computing so-called Principal Components. These computations require finding an eigenvector decomposition, and again require complex iterative algorithms. While these more complex algorithms are well-known and understood, they are nevertheless vastly more complex than the methods of the present invention.

3. The methods of the present invention are more robust than other methods. Methods of the prior art typically depend on finding the solution to a system of linear equations under constraints, or in solving for linear invariants such as

eigenvectors. It is often hard to predict how the stability of these methods depends on the data and assumptions. Due to the straightforward nature of the methods of the present invention, unstable solutions are easily avoided.

5
4. The methods of the present invention are amenable to use in a system where partial solutions can be used. For example, in considering a large number of candidate models for identification, many models can be discarded as soon as the error is known to be sufficiently large. Using the methods of the present invention, it is possible to interrupt computation of the error or quality term at numerous stages, and resume it only if needed at later stages. In the methods of the prior art, it is difficult or impossible to do this without incurring the full cost of the most expensive part of the method.

10
15
20
5. The methods of the present invention are adaptable to unknown models of reflectance. Because light sources need not be computed, only the effect of lighting on the material being imaged is measured and assumed consistent. Even with specularities (shininess), the methods of the present invention allow for good results over a reasonable range of material properties.

25
6. The methods of the present invention are robust against errors in capturing reflectance, particularly if lighting effects are present in the captured reflectance data.

30
Since the methods of the present invention are extremely fast and do not require a great amount of computational effort, they may also be used in a system to winnow the candidate models that will be examined by more costly (and slower) methods by a factor of 10-100. These other methods are capable of a more precise identification, but are too costly for a very large database of

candidate models.

While there has been shown and described what is considered to be preferred embodiments of the invention, it will, of course, be understood that various modifications and changes in form or detail could readily be made without departing from the spirit of the invention. It is therefore intended that the invention be not limited to the exact forms described and illustrated, but should be constructed to cover all modifications that may fall within the scope of the appended claims.

WHAT IS CLAIMED IS:

1 1. A method for matching a two dimensional image to one of
2 a plurality of three dimensional candidate models, the method
3 comprising the steps of:

4
5 determining the position and orientation of the two
6 dimensional image;

7
8 for each three dimensional candidate model, computing a
9 histogram-like table having a computed brightness coefficient for
10 each surface normal of the three dimensional candidate model
11 based on the corresponding value in the two dimensional image,
12 each brightness coefficient depending only on the corresponding
13 surface normal;

14
15 successively rendering each three dimensional candidate
16 model in the determined position and orientation using the
17 surface normals in conjunction with the corresponding computed
18 brightness histogram-like table; and

19
20 comparing the two dimensional image with each of the
21 rendered three dimensional candidate models.

22
23 2. The method of claim 1, further comprising the steps of
24 determining if the two dimensional image matches one of the
25 plurality of three dimensional models if the comparison between
26 the two dimensional image and the rendered three dimensional
27 model is within a predetermined allowable error.

28
29 3. The method of claim 1, wherein the comparison step
30 results in a ranking for each of the rendered three dimensional
31 models which indicates the likelihood that the corresponding
32 three dimensional model matches the two dimensional image.

1 4. The method of claim 3, further comprising the step of
2 choosing the three dimensional model corresponding to the
3 rendered three dimensional model having the highest ranking as
4 the three dimensional model which matches the two dimensional
5 image.

1 5. The method of claim 1, further comprising the step of
2 reading the three dimensional models from a database containing
3 the plurality of three dimensional models.

1 6. The method of claim 1, wherein the comparing step is
2 done simultaneously with the rendering step for each rendered
3 three dimensional model.

1 7. The method of claim 1, wherein the computing step sets
2 a minimum threshold of reflectance to be used in the computation
3 of the histogram-like table.

1 8. The method of claim 1, wherein the computing step sets
2 a minimum and/or maximum threshold of intensity to be used in the
3 computation of the histogram-like table.

1 9. The method of claim 1, wherein the two dimensional
2 image is a color image and wherein the a brightness coefficient
3 is computed for each color component of the color image.

1 10. The method of claim 1, further comprising the step of
2 smoothing the lightsphere histogram-like table to improve the
3 rendering of each three dimensional model and the performance of
4 the method.

1 11. The method of claim 10, wherein the smoothing step is
2 done by a convolution with a small box function or other small

3 smoothing kernel.

1 12. A method for matching a two dimensional image to one of
2 a plurality of three dimensional candidate models, the method
3 comprising the steps of:

4
5 determining the position and orientation of the two
6 dimensional image;

7
8 for each three dimensional candidate model, computing a
9 histogram-like table having a computed brightness coefficient for
10 each surface normal of the three dimensional candidate model
11 based on the corresponding value in the two dimensional image,
12 each brightness coefficient depending only on the corresponding
13 surface normal;

14
15 computing the variance of the brightness coefficients
16 that are used to create each bucket of the histogram-like table,
17 a bucket being a set of similar normal values that are lumped
18 together to form a single argument value in the histogram-like
19 table;

20
21 computing the sum of the variances of the histogram-
22 like table buckets; and

23
24 ranking the three dimensional candidate models using
25 its computed sum as an error function, the ranking indicating the
26 likelihood that the corresponding three dimensional model matches
27 the two dimensional image.

1 13. The method of claim 12, further comprising the step of
2 choosing the three dimensional model corresponding to the three
3 dimensional model having the highest ranking as the three
4 dimensional model which matches the two dimensional image.

1 14. The method of claim 12, further comprising the step of
2 reading the three dimensional models from a database containing
3 the plurality of three dimensional models.

1 15. The method of claim 12, wherein the computing step sets
2 a minimum threshold of reflectance to be used in the computation
3 of the histogram-like table.

1 16. The method of claim 12, wherein the computing step sets
2 a minimum and/or maximum threshold of intensity to be used in the
3 computation of the histogram-like table.

1 17. The method of claim 12, wherein the two dimensional
2 image is a color image and wherein the a brightness coefficient
3 is computed for each color component of the color image.

1 18. The method of claim 12, wherein the sum of the
2 variances are weighted by the number of samples in each bucket.

1 19. A method for matching a two dimensional image to one of
2 a plurality of three dimensional candidate models, the method
3 comprising the steps of:

4 determining the position and orientation of an object
5 giving rise to the two dimensional image;

6
7
8 computing a representation of lighting effects that
9 allows the lighting that gave rise to the two dimensional image
10 to be used to render a realistic image of a three dimensional
11 candidate model;

12
13 successively rendering each three dimensional candidate
14 model in the determined position and orientation using the

15 surface normals in conjunction with the corresponding computed
16 representation of lighting effects; and

17
18 comparing the two dimensional image with each of the
19 rendered three dimensional candidate models.

1 20. A method for matching a two dimensional image to one of
2 a plurality of three dimensional candidate models, the method
3 comprising the steps of:

4
5 determining the position and orientation of an object
6 giving rise to the two dimensional image;

7
8 computing a representation of lighting effects based on
9 the plurality of three dimensional candidate models and the two
10 dimensional image that allows evaluation of the likelihood that a
11 particular three dimensional candidate model gave rise to a
particular two dimensional image; and

12
13 choosing the most likely three dimensional candidate
14 model to have generated the query based on the computed
15 representation of lighting effects.
16

17 21. A computer program product embodied in a computer-
18 readable medium for matching a two dimensional image to one of a
19 plurality of three dimensional candidate models stored in a
20 database, the computer program product comprising:

21
22 computer readable program code means for determining
23 the position and orientation of the two dimensional image;

24
25 computer readable program code means for computing a
26 histogram-like table each three dimensional candidate model,
27 having a computed brightness coefficient for each surface normal

12 of the three dimensional candidate model based on the
13 corresponding value in the two dimensional image, each brightness
14 coefficient depending only on the corresponding surface normal;

15
16 computer readable program code means for successively
17 rendering each three dimensional candidate model in the
18 determined position and orientation using the surface normals in
19 conjunction with the corresponding computed brightness histogram-
20 like table; and

21
22 computer readable program code means for comparing the
23 two dimensional image with each of the rendered three dimensional
24 candidate models.

25 22. A program storage device readable by machine, tangibly
26 embodying a program of instructions executable by the machine to
27 perform method steps for matching a two dimensional image to one
28 of a plurality of three dimensional candidate models, the method
29 comprising the steps of:

30 determining the position and orientation of the two
31 dimensional image;

32 for each three dimensional candidate model, computing a
33 histogram-like table having a computed brightness coefficient for
34 each surface normal of the three dimensional candidate model
35 based on the corresponding value in the two dimensional image,
36 each brightness coefficient depending only on the corresponding
37 surface normal;

38 successively rendering each three dimensional candidate
39 model in the determined position and orientation using the
40 surface normals in conjunction with the corresponding computed
brightness histogram-like table; and

21 comparing the two dimensional image with each of the
22 rendered three dimensional candidate models.

1 23. A computer program product embodied in a computer-
2 readable medium for matching a two dimensional image to one of a
3 plurality of three dimensional candidate models stored in a
4 database, the computer program product comprising:

5
6 computer readable program code means for determining
7 the position and orientation of the two dimensional image;

8
9 computer readable program code means for computing a
10 histogram-like table for each three dimensional candidate model,
11 having a computed brightness coefficient for each surface normal
12 of the three dimensional candidate model based on the
13 corresponding value in the two dimensional image, each brightness
14 coefficient depending only on the corresponding surface normal;

15
16 computer readable program code means for computing the
17 variance of the brightness coefficients that are used to create
18 each bucket of the histogram-like table, a bucket being a set of
19 similar normal values that are lumped together to form a single
20 argument value in the histogram-like table;

21
22 computer readable program code means for computing the
23 sum of the variances of the histogram-like table buckets; and

24
25 computer readable program code means for ranking the
26 three dimensional candidate models using its computed sum as an
27 error function, the ranking indicating the likelihood that the
28 corresponding three dimensional model matches the two dimensional
29 image.

1 24. A program storage device readable by machine, tangibly

2 embodying a program of instructions executable by the machine to
3 perform method steps for matching a two dimensional image to one
4 of a plurality of three dimensional candidate models, the method
5 comprising the steps of:

6
7 determining the position and orientation of the two
8 dimensional image;

9
10 for each three dimensional candidate model, computing a
11 histogram-like table having a computed brightness coefficient for
12 each surface normal of the three dimensional candidate model
13 based on the corresponding value in the two dimensional image,
14 each brightness coefficient depending only on the corresponding
15 surface normal;

16
17 computing the variance of the brightness coefficients
18 that are used to create each bucket of the histogram-like table,
19 a bucket being a set of similar normal values that are lumped
20 together to form a single argument value in the histogram-like
21 table;

22
23 computing the sum of the variances of the histogram-
24 like table buckets; and

25
26 ranking the three dimensional models using its computed
27 sum as an error function, the ranking indicating the likelihood
28 that the corresponding three dimensional candidate model matches
29 the two dimensional image.

1 25. A computer program product embodied in a computer-
2 readable medium for matching a two dimensional image to one of a
3 plurality of three dimensional candidate models stored in a
4 database, the computer program product comprising:

6 computer readable program code means for determining
7 the position and orientation of an object giving rise to the two
8 dimensional image;

10 computer readable program code means for computing a
11 representation of lighting effects that allows the lighting that
12 gave rise to the two dimensional image to be used to render a
13 realistic image of a three dimensional candidate model;

15 computer readable program code means for successively
16 rendering each three dimensional candidate model in the
17 determined position and orientation using the surface normals in
18 conjunction with the corresponding computed representation of
19 lighting effects; and

21 computer readable program code means for comparing the
22 two dimensional image with each of the rendered three dimensional
23 candidate models.

25 26. A program storage device readable by machine, tangibly
26 embodying a program of instructions executable by the machine to
27 perform method steps for matching a two dimensional image to one
28 of a plurality of three dimensional candidate models, the method
29 comprising the steps of:

31 determining the position and orientation of an object
32 giving rise to the two dimensional image;

34 computing a representation of lighting effects that
35 allows the lighting that gave rise to the two dimensional image
36 to be used to render a realistic image of a three dimensional
37 candidate model;

39 successively rendering each three dimensional candidate

16 model in the determined position and orientation using the
17 surface normals in conjunction with the corresponding computed
18 representation of lighting effects; and

19
20 comparing the two dimensional image with each of the
21 rendered three dimensional candidate models.

1 27. A computer program product embodied in a computer-
2 readable medium for matching a two dimensional image to one of a
3 plurality of three dimensional candidate models stored in a
4 database, the computer program product comprising:

5
6 computer readable program code means for determining
7 the position and orientation of an object giving rise to the two
8 dimensional image;

9
10 computer readable program code means for computing a
11 representation of lighting effects based on the plurality of
12 three dimensional candidate models and the two dimensional image
13 that allows evaluation of the likelihood that a particular three
14 dimensional candidate model gave rise to a particular two
15 dimensional image; and

16
17 computer readable program code means for choosing the
18 most likely three dimensional candidate model to have generated
19 the query based on the computed representation of lighting
20 effects.

1 28. A program storage device readable by machine, tangibly
2 embodying a program of instructions executable by the machine to
3 perform method steps for matching a two dimensional image to one
4 of a plurality of three dimensional candidate models, the method
5 comprising the steps of:
6

7 determining the position and orientation of an object
8 giving rise to the two dimensional image;

9
10 computer readable program code means for computing a
11 representation of lighting effects based on the plurality of
12 three dimensional candidate models and the two dimensional image
13 that allows evaluation of the likelihood that a particular three
14 dimensional candidate model gave rise to a particular two
15 dimensional image; and

16
17 computer readable program code means for choosing the
18 most likely three dimensional candidate model to have generated
19 the query based on the computed representation of lighting
20 effects.

**METHOD FOR MATCHING A TWO DIMENSIONAL IMAGE TO ONE OF A PLURALITY
OF THREE DIMENSIONAL CANDIDATE MODELS CONTAINED IN A DATABASE**

ABSTRACT OF THE DISCLOSURE

25

30

40

45

50

A method for matching a two dimensional image to one of a plurality of three dimensional candidate models. The method including the steps of: determining the position and orientation of the two dimensional image; for each three dimensional model, computing a histogram-like table having a computed brightness coefficient for each surface normal of the model based on the corresponding value in the two dimensional image, each brightness coefficient depending only on the corresponding surface normal; and either successively rendering each three dimensional model in the determined position and orientation using the surface normals in conjunction with the corresponding computed brightness histogram-like table; and comparing the two dimensional image with each of the rendered three dimensional models or computing the variance of the brightness coefficients that are used to create each bucket of the histogram-like table, a bucket being a set of similar normal values that are lumped together to form a single argument value in the histogram-like table; computing the sum of the variances of the histogram-like table buckets; and ranking the three dimensional models using its computed sum as an error function, the ranking indicating the likelihood that the corresponding three dimensional model matches the two dimensional image. Also provided is a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform the method steps of the present invention and a computer program product embodied in a computer-readable medium for carrying out the methods of the present invention.

Fig. 1


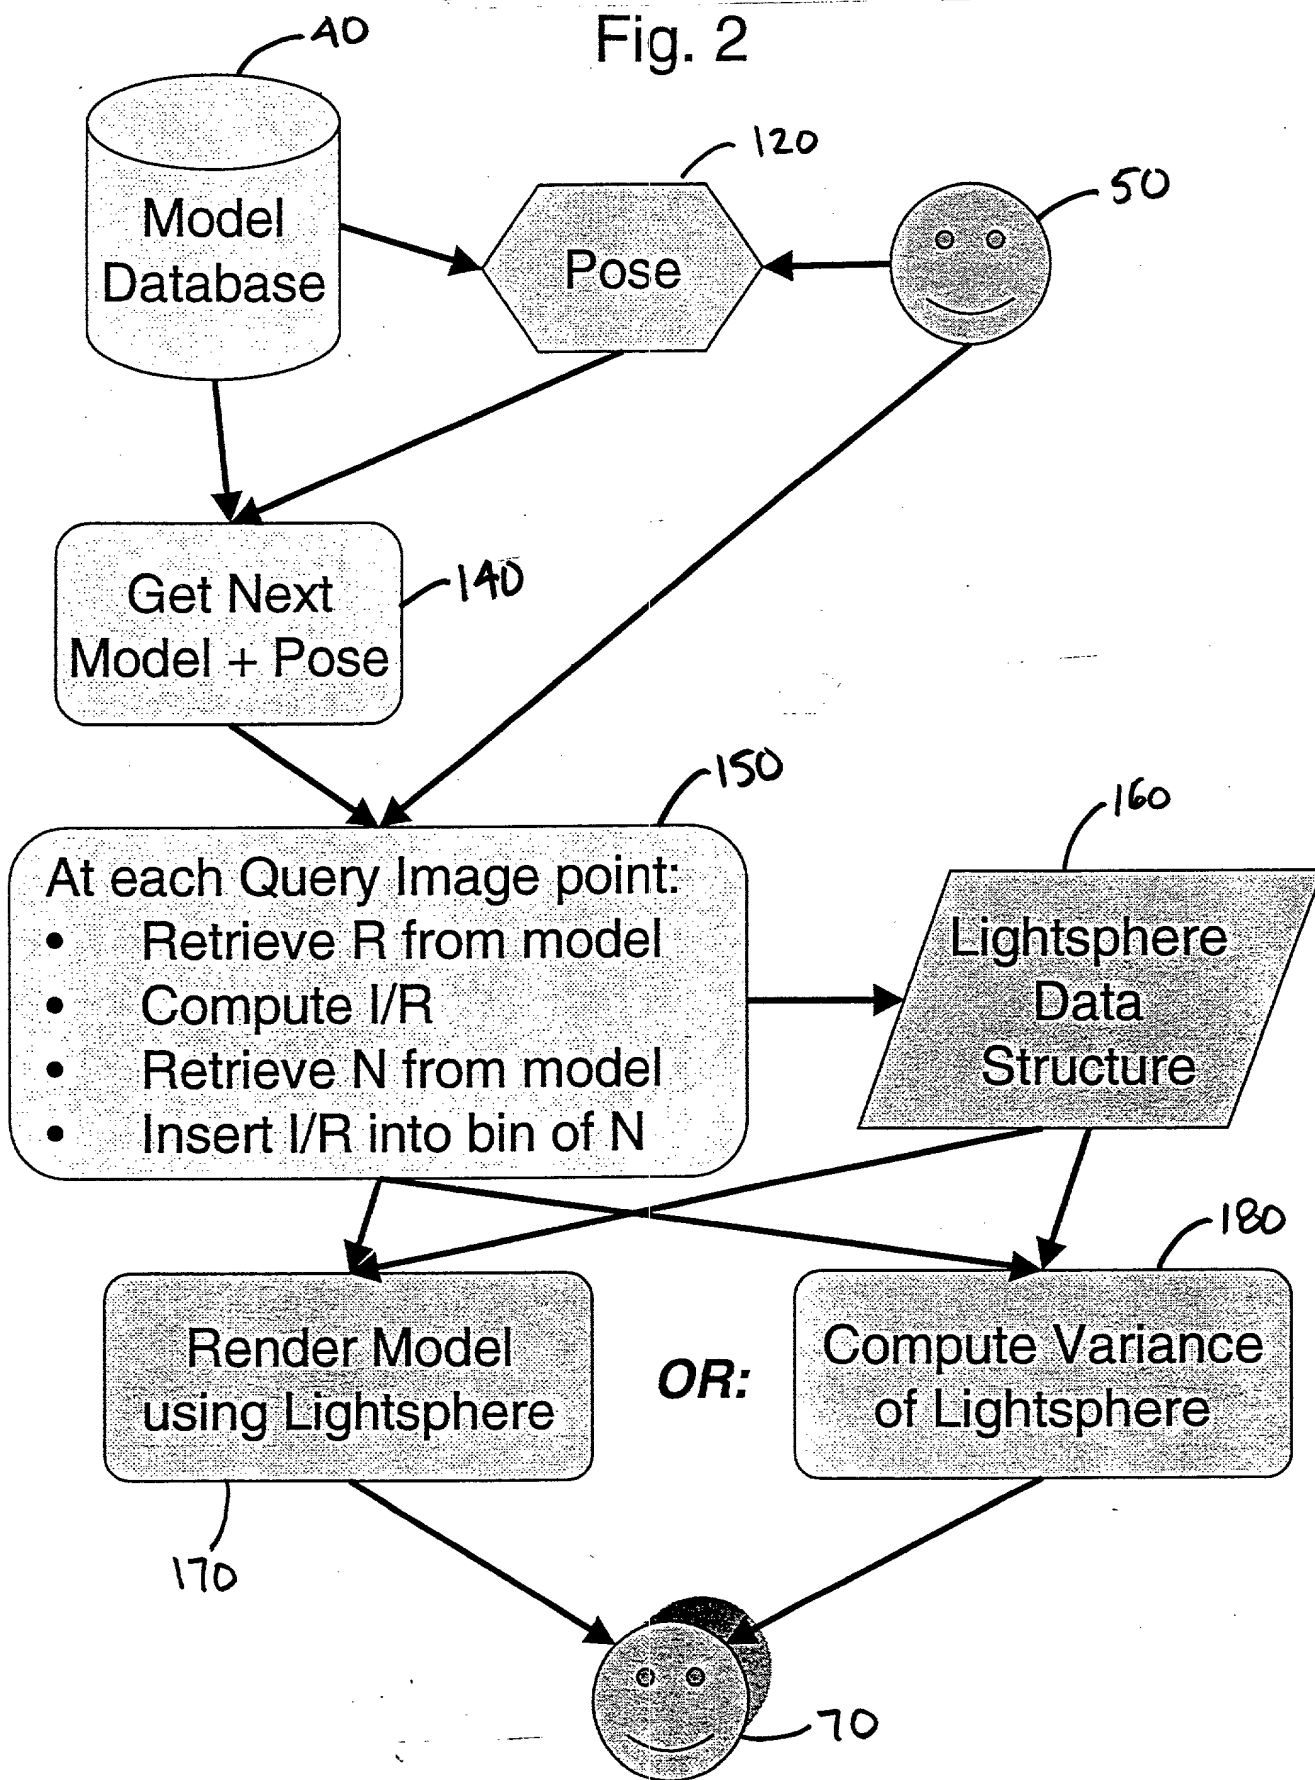



Fig. 2



[illegible]

Docket No.
13414 (NECI1083)

Declaration and Power of Attorney For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD FOR MATCHING A TWO DIMENSIONAL IMAGE TO ONE OF A PLURALITY OF THREE DIMENSIONAL CANDIDATE MODELS CONTAINED IN A DATABASE

the specification of which

(check one)

☒ is attached hereto.

☐ was filed on _____ as United States Application No. or PCT International Application Number _____ and was amended on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)

Priority Not Claimed

(Number)

(Country)

(Day/Month/Year Filed)

☐

(Number)

(Country)

(Day/Month/Year Filed)

☐

(Number)

(Country)

(Day/Month/Year Filed)

☐

I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, C. F. R., Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. *(list name and registration number)*

Stephen D. Murphy; Reg. No. 22,002

William C. Roch; Reg. No. 24,972

Frank S. DiGiglio; Reg. No. 31,346

Paul J. Esatto, Jr.; Reg. No. 30,749


Mark J. Cohen; Reg. No. 32,211


Philip J. Feig; Reg. No. 27,328

Andrew G. Isztwan; Reg. No. 40,028

Send Correspondence to: **Paul J. Esatto, Jr., Esq.**
Scully, Scott, Murphy & Presser
400 Garden City Plaza
Garden City, New York 11530

Direct Telephone Calls to: *(name and telephone number)*
Paul J. Esatto, Jr. Esq. at (516) 742-4343

| | |
|---|--------------------------|
| Full name of sole or first inventor Sebastien Roy | |
| Sole or first inventor's signature  | Date 3-27-2000 |
| Residence Montreal, Canada | |
| Citizenship Canada | |
| Post Office Address 4126 ave. Laval | |
| Montreal (Quebec), H2W 2J3, Canada | |

| | |
|--|------------------------|
| Full name of second inventor, if any Albert Peter Blicher | |
| Second inventor's signature  | Date 3/29/00 |
| Residence East Windsor, NJ | |
| Citizenship U.S.A. | |
| Post Office Address 54 Spruce Road | |
| East Windsor, NJ 08520 | |